



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 905 643 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
31.03.1999 Bulletin 1999/13

(51) Int Cl.⁶: G06K 9/00

(21) Application number: 98307406.3

(22) Date of filing: 14.09.1998

BEST AVAILABLE COPY

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Syeda-Mahmood, Tanveer F.
Rochester, New York 14625 (US)

(74) Representative: Skone James, Robert Edmund
GILL JENNINGS & EVERY
Broadgate House
7 Eldon Street
London EC2M 7LH (GB)

(30) Priority: 29.09.1997 US 939894

(71) Applicant: XEROX CORPORATION
Rochester, New York 14644 (US)

(54) Method and system for recognizing handwritten words

(57) A method and system of recognizing handwritten words in scanned documents, wherein by processing a document containing handwriting, features for word localization are extracted from handwritten words contained in said document through basis points taken from a single curve of text lines. The method is independent of page orientation, and does not assume that the individual lines of handwritten text are parallel, and the method does not require that word regions be aligned with text line orientation wherein intra-word statistics are derived from sample pages rather than using a fixed threshold. The method has applications in digital libraries, handwriting tokenization, document management and OCR systems.

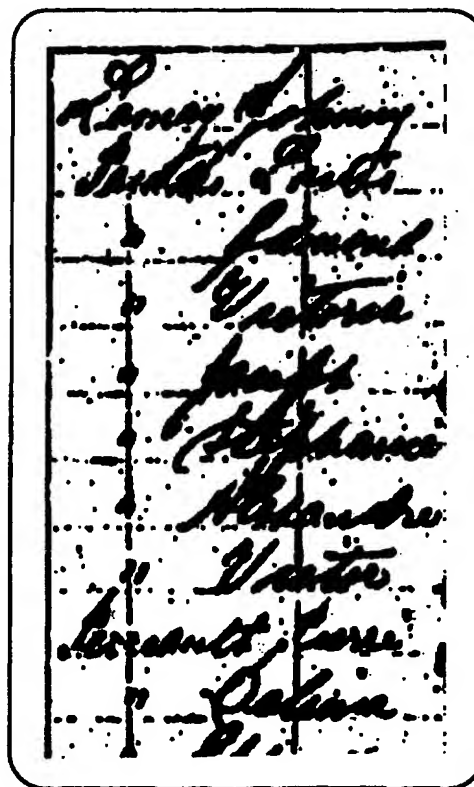


FIG. 1

Description

[0001] This invention is related to handwriting recognition and indexing and, more particularly, to a method for grouping text segments into handwritten words for purposes of indexing such documents based on word queries.

[0002] The ability to detect and recognize handwritten words in handwritten documents is important for several applications. While the strategic importance of such a capability in current commercial handwriting recognition products is clear, its use in applications such as digital libraries and document management cannot be ignored. With digital libraries, for example, there is a major concern over the preservation and electronic conversion of historical paper documents. Often, these documents are handwritten and in calligraphic styles, as in a sample of a church record used in genealogy studies illustrated in Figure 1. An important aspect of the use of electronic versions of such documents is their access based on word queries. Handwritten keyword extraction and indexing can also be a valuable capability for document management, in handling a variety of irregular paper documents such as handwritten notes, marks on engineering drawings, memos and legacy documents.

[0003] While an OCR algorithm can be used to extract text keywords for index creation of scanned printed text documents, such a process is not yet an option for handwritten documents due to a lack of robust handwriting recognition algorithms. One of the difficulties is due to the fact that the same word could be written differently at different locations in a document even when the document is written by a single author. In cursive script, this often means that a word is written as a collection of word segments separated by intra-word separations that are characteristic of the author. Figures 2A-C illustrate this situation, where the word "database" is written by the same author differently in the various instances it occurs. Further, the different word instances could exhibit different amounts of global skew, because lines of handwritten text are often not parallel as in printed text. This latter fact makes the detection of lines of handwritten text a further difficulty during recognition.

[0004] The present method of grouping handwritten words was motivated by an application that required image indexing of old calligraphic handwritten church record documents for purposes of tracing genealogy. These documents were written against a tabular background, as shown in Figure 1. On being given a query about a person's name, the task was to locate the relevant records. While the formulation of query word patterns for these documents is an interesting problem, for the purposes of this disclosure relevant problem is that of matching handwritten words after they have been formulated by a user -- perhaps by a training process that generates such pattern queries from actual typed text queries, or perhaps such queries are derived from the handwritten document itself.

[0005] A method of localizing handwritten word patterns in documents exploiting a data structure, called the image hash table, to succinctly represent feature information needed to localize any word without a detailed search of the document, is presented in a co-pending patent application, Serial No. 08/878,512, filed June 1, 1997, by the assignee in this case. The use of an image hash table to localize objects draws upon ideas of geometric hashing that has been used in the past for identification of objects in pre-segmented image regions. These concepts are discussed in articles by Y. Lamdan and H.J. Wolfson entitled "Geometric hashing: A general and efficient model-based recognition scheme", Proceeding of the International Conference on Computer Vision, pages 218-249, 1988, and "Transformation invariant indexing", Geometric Invariants in Computer Vision, MIT Press, pages 334-352, 1992. More work has been done in extending the basic geometric hashing scheme for use with line features as described in an article by F.C.D. Tsai entitled "Geometric hashing with line features" Pattern Recognition, Vol. 27, No. 3, pages 377-389, 1994. An extensive analysis of the geometric hashing scheme has been done in an article by W.E.L. Grimson and D. Huttenlocher entitled "On the sensitivity of geometric hashing", Proceedings International Conference on Computer Vision, pages 334-339, 1990. Finding good geometric hash functions has also been explored in an article by G. Bebis, M. Georgiopolous and N. Lobo entitled "Learning geometric hashing functions for model-based object recognition", Proceedings International Conference on Computer Vision, pages 543-548, 1995, and an extension of geometric hashing using the concept of rehashing the hash table has been discussed in an article by I. Rigoustos and R. Hummel "Massively parallel model matching: Geometric hashing on the connection machine", IEEE Computer, pages 33-41, February 1992.

[0006] All the prior work has used the geometric hashing technique for purposes of model indexing in object recognition where the task is to determine which of the models in a library of models is present in the indicated region in the image. The localization of handwritten words in unsegmented handwritten documents is an instance of image indexing (rather than model indexing) for which no prior work on using geometric hashing exists. Work that uses a serial search of the images for localizing handwritten words as described in an article by R. Manmatha, C. Han and E. Riseman, entitled "Word spotting: A new approach to indexing handwriting", Proceedings IEEE Computer Vision and Pattern Recognition Conference, pages 631-637, 1996, only begins to address the need.

[0007] U.S. Patent No. 5,640,466 issued to Huttenlocher et al. on June 17, 1997, entitled "Method of Deriving Word-shapes for Subsequent Comparison", describes a method for reducing an image of a character or word string to one or more one dimensional signals, including steps of determining page orientation, isolating character strings from adjacent character strings, establishing a set of references with respect to which measurement about the character string may be made, and driving a plurality of measurements with respect to the references in terms of a single variable

signal, from which information about the symbol string may be derived.

[0008] Localization or indexing of a specific word in the document is done by indexing the hash table with information derived from the word in such a manner that the prominent hits in the table directly indicate candidate locations of the word in the document, thus avoiding a detailed search. This method accounts for changes in appearance of the handwritten word in terms of orientation, skew, and intra-word separation that represent the way a single author may write the same word at different instances. More specifically, localizing any word in the image hash table is done by indexing the hash table with features computed from the word pattern. The top hits in the table are candidate locations most likely to contain the word. Such an indexing automatically gives pose information which is then used to project the word at the indicated location and then verify it. Verification then involves determining the extent of match between the underlying word and the projected word. The generation and indexing of image hash tables takes into account the changes in appearance of the word under 2D affine transforms, changes in the orientation of the lines of text, overall document skew, changes in word appearance due to occlusions, noise, or intra-word handwriting variations made by a single author.

[0009] Generally, localization and detection of handwritten words involves four stages: (1) Pre-processing where features for word localization are extracted; (2) Image hash table construction; (3) Indexing where query word features are used to look up hash table for candidate locations; and (4) Verification, where the query word is projected and registered with the underlying word at the candidate locations. The focus of the present disclosure is on stage (1) of this processing, namely, in the stage where features for word localization are generated. Therefore, a feature of the present invention is in the ability to recognize and generate handwritten word regions for purposes of feature generation used ultimately for handwritten word indexing.

[0010] In accordance with one aspect of the present invention, a method of recognizing handwritten words in scanned documents comprises processing a document containing handwriting wherein features for word localization are extracted from handwritten words contained in said document through basis points taken from a single curve of text lines, and wherein affine coordinates are computed for all features on all curves in a curve group; storing said features in a memory; accessing said features from memory for comparison to handwritten words in a scanned document to recognize words within said scanned document.

[0011] The invention is a method of grouping text segments to generate handwritten words for the recognition and indexing of documents. An ability to accomplishing handwritten word indexing not only extends the capability of current document management systems by allowing handwritten documents to be treated in a uniform manner with printed text documents but can also be the basis for compressing such documents by handwritten word tokenization.

[0012] As discussed earlier, localization and detection of handwritten words generally involves four stages: (1) Pre-processing where features for word localization are extracted; (2) Image hash table construction; (3) Indexing where query word features are used to look up hash table for candidate locations; and (4) Verification, where the query word is projected and registered with the underlying word at the candidate locations. The focus of the present disclosure is on stage (1), the pre-processing of handwritten words, namely, in the stage where features for word localization are generated. Specifically, this stage presents a method for grouping text segments into handwritten words by doing the following processing stages: (1) Connected region generation; (2) Region feature extraction; (3) Orientation histogram computation; (4) Selective Hough transform generation; (5) Handwritten text line detection; (6) Along-line inter-region distance computation; (7) Intra-word separation determination; (8) Curve and corner feature extraction from regions; and finally, (9) Intra-word text segment grouping.

[0013] As for a system, a microprocessor can be programmed to generate words in handwritten documents by processing a scanned document containing handwriting to extract and group features from handwritten words contained in the document.

[0014] The accompanying drawings, which are incorporated in and form part of the specification, illustrate an embodiment of the present invention and, together with the description, serve to better explain the operation features, and advantages of the invention. It should be understood, however, that the invention is not limited to the precise arrangements and instrumentalities shown.

Figure 1 illustrates a scanned image of a sample handwritten document;

Figure 2A illustrates a second sample handwritten document image;

Figure 2B illustrates a handwritten query word "database" within the sample handwritten document of Figure 2A;

Figure 2C illustrate the subject query word "database" projected at candidate locations within the scanned handwritten document image of Figure 2A;

Figure 3 illustrates a block diagram of the processing modules of the invention involved in hand written word recognition;

Figure 4 illustrates a an orientation histogram of text regions in the image of Figure 5A;

Figure 5 illustrates curves in the handwritten sample document of Figure 2A;

Figure 6 illustrates lines of text groupings to peaks as determined by a selective Hough transform;

Figure 7 illustrates a peak of at least 2 separations from the histogram;

Figure 8 illustrates the separations corresponding to the peak at the lowest separation value used as an estimate of intra-word separation in the algorithm;

Figure 9 illustrated a block diagram of system modules implementing the invention for handwritten word recognition and also engaged in handwritten word group generation and indexing;

Figure 10 is a further illustration of a system block diagram of modules engaged in query localization by image indexing of hash tables using the pre-processing performed by the current invention;

Figure 11A illustrates the results of preprocessing and feature extraction of the image of Figure 1;

Figure 11B illustrates a query pattern consisting of a single curve extracted image of Figure 11A;

Figure 12 illustrates a histogram of hashing based coordinates for Figure 11B; and,

Figure 13 illustrates Hashing results for the histogram of Figure 12.

[0015] This invention is about a method of grouping text segments forming part of single handwritten words in a document. This is done primarily to enable handwriting localization under changes in word appearance using an image hash table data structure, or similar devices known in the character recognition art, populated with features derived from text regions. The grouping of text segments into handwritten words requires the knowledge of intra-word separation between text segments that lie along a line of text. Unlike in printed text, deducing lines of text in handwritten document usually is difficult because handwritten text words are often not written on a straight line. Furthermore, consecutive lines of text may not be parallel as in printed text. Finally, an author may vary the inter-word and intra-word spacing while writing so that different instances of the same word may show writing differences. This makes the task of determining which text segments belong to a word difficult.

[0016] The method of text lines detection that is disclosed herein is independent of page orientation, and does not assume that the individual lines of handwritten text are parallel. Furthermore, it does not require that all word regions be aligned with the text line orientation. Finally, it derives intra-word statistics from the sample page itself, rather than using a fixed threshold.

[0017] Referring to Figure 3, the components for generating handwritten word regions in documents are illustrated. In the pre-processing step of the invention, original documents are obtained by scanning handwritten pages at high resolution (typically 200dpi or higher) are used.

[0018] In the first step, connected regions of text in the scanned document are formed in the connected region module 21. Although several methods of finding connected components exist, the following algorithm is used to determine the connected components regions in bitmaps:

1. Record run lengths of "on" pixels (assuming white background) per image pixel row using low[i], high[i] arrays that maintain the start and end points of the run lengths.
2. Initially put all runlengths in separate groups denoted by C__i for runlength i.
3. For all end point pixels (k,l) in low[i] and high[i] arrays, do the following steps:

Step A: Find the number of "on" neighboring pixels (k',l') and their associated run lengths, and

Step B: Merge the given runlength with the neighboring runlength identified above. This is recorded by having all merged runlength having the same group identification.

[0019] The above algorithm can be efficiently implemented using a data structure called the union-find data structure as described in a book by Cormen, Leiserson and Rivest entitled "Introduction to algorithms", MIT Press, 1994, to run in time linear in the number of runlengths in the image.

[0020] In the next stage of processing, region features such as the centroid and dominant orientation are noted in the region feature extraction module 22. The centroid of the region is derived from the first order moments of the region, while the orientation of the region is determined by the direction of the moment of inertia axis of the region. The formula for finding the moment of inertia axis is given in Chapter 3 of the book entitled "Robot Vision" by B. K.P. Horn, MIT Press, 1986, and is re-produced here for convenience.

[0021] The orientation of the moment of inertia axis is given by the orientation of the line joining the origin '(0,0) with the point

$$\left(\sqrt{\sqrt{(a-c)^2 + b^2} - (c-a)}, \sqrt{\sqrt{(a-c)^2 + b^2} - (a-c)} \right).$$

Where a, b, and c, are the second order moments given by:

$$a = \iint x^2 b(x,y) dx dy, \quad b = 2 \iint xy b(x,y) dx dy, \quad c = \iint y^2 b(x,y) dx dy, \quad \text{where } b(x,y) = 1$$

when the pixels are on the region and 0 otherwise, as determined from the connected component generation.

[0022] A histogram of orientations is generated in orientation histogram computation module 23. Peaks in the histogram are automatically selected to represent major word orientations in the image. For each of the dominant orientations selected a line of the specified orientation is drawn through the centroids of each of the regions. The selective Hough transform module 24 performs a selective Hough transform to determine groups of such lines. The Hough transform described in a book by D. Ballard and C. Brown entitled "Computer Vision", Prentice-Hall, Chapter 4, pages 123-124, 1982, was used to record this information. The resulting data structure, called the Hough Transform Table, is a two-dimensional array that records the number of points (centroids of region here) that lie along or lie close to a line of specified orientation and position. The highest valued entries in this table are taken to correspond to candidate lines of text in the handwritten text line detection module 25. The regions whose centroids contribute to the peak table entries are noted. These word segment regions thus are taken to form the lines of text in the handwritten document image.

[0023] Once the lines of text, and hence the word segments that lie along a lines of text, are determined, the intra-word separation is estimated in the Intra-word separation determination module 26. For each line of text determined above, the boundaries of the word segment regions lying on the line are used to determine two extremal points per region; that is, all the boundary points of a region are projected onto the line, and the beginning and end points noted. A projection of a given point onto a line is the point of intersection of a perpendicular line through the given point with the given line. All such projections are now sorted in an increasing order along the line, using a conventional sorting algorithm. Distances between the end point of a region and the beginning point of another are noted to represent separations between word segments. These distances are recorded for all lines of text. A histogram of such distances is generated. For most handwritten documents such a histogram shows at least two distinct peaks. The peak at the lowest separation distance is noted as intra-word separation.

[0024] Using the estimated intra-word separation, the text segmented belonging to the words are generated by extracting curves from the text regions. The curve feature extraction module 27 proceeds by determining the boundary points on the connected component regions as those points that have at least one "off" neighbor. A cyclic trace of such boundary pixels is used to yield curves representing the boundaries of the connected component regions. The curves are smoothed using a conventional line-segment approximation algorithm. Other methods of curve tracing can be used without significantly affecting the claims in this invention.

[0025] The pre-processing step of curve extraction can be applied uniformly to a document image or to a query word represented as an image pattern, and takes time linear in the size of the image.

[0026] Using the intra-word separation, and the curves belonging to text regions, the intra-word text segment grouping module 28 assembles curve groups by grouping word segment regions that are separated along the line of text orientation by a distance within a certain bound of the intra-word separation determined above. The grouping of curves separated by intra-word separation (+/- a chosen threshold) is done using the union-find data structure mentioned earlier.

Example

[0027] Figure 2A shows a scanned handwritten document in which the word database appears segmented differently into text segments in the instances in which it occurs. Figure 4 shows the orientation histogram of the text regions, clearly showing a peak at nearly horizontal text line orientation. By taking a selective Hough transform along the selected peak orientation, and grouping text regions on lines corresponding to peaks in the Hough transform as described above and shown in Figure 5, the lines of text determined by the Hough transform are shown in Figure 6. Here the text segments belonging to a line are rendered in identical line thickness. Next, Figures 7 and 8 illustrate the text segment grouping for forming handwritten words in a document. The histogram of separations of Figure 7 show a peak at least two separations. The separation corresponding to the peak at the lowest separation value is used as an estimate of intra-word separation in the algorithm as shown in Figure 8.

Use of grouped text segments in word localization

[0028] This section describes the use of the handwritten word for purposes of query localization as also reported in the co-pending patent application identified in the background. From the curves of the text regions grouped as words, corner features are derived as those points where significant curvature deviation occurs. That is, where the angle between two incident lines is greater than a specified threshold. Note that since the images are assumed to be scanned at high resolution, the lines are thick enough so that junctions are also manifested as corners in such images. Corner

features on a curve are chosen as the basic unit for localization using the rationale that although not all curves come from single words, especially in the presence of occlusions and noise, features generated from within a curve are more likely to point to a single image location than an arbitrary triple of features chosen randomly across the image.

[0029] Using the features derived above, a data structure called an image hash table is developed within the Hash Table Construction Module 4 in Figure 9 and is used to succinctly represent information in the position of features in curves in curve groups in a manner that helps locate a query handwritten word. To understand the idea of an image hash table, suppose for the sake of simplicity, each curve group consists of a single curve. Suppose the task is to locate a given query curve in an image consisting of this curve among others. Consider three consecutive non-collinear feature points (O, P₁, P₂) on the given query curve. Then it is well-known that the coordinates of any other point P of the curve can be expressed in terms of the coordinates of points (O, P₁, P₂) (called basis triples) as:

$$OP = \alpha OP_1 + \beta OP_2$$

[0030] The coordinates (α, β) are called affine coordinates and they are invariant to affine transformations. Thus, if the given curve appears in the image skewed, or rotated, the corresponding points on the transformed image curve will have the same coordinates with respect to the transformed basis triples in the transformed image curve. Thus, one way to check if a curve at an image location matches a given curve is to see if enough feature points on the image curve have the same affine coordinates with respect to some image basis triple (O', P'₁, P'₂) on the image curve. In this case, it can also be inferred that the basis triples on the image curve and the given (query) curve correspond. From such a correspondence, the pose information can be derived as an affine transform:

$$(A, T) = \begin{pmatrix} a & b \\ c & d \end{pmatrix} + \begin{pmatrix} T_1 \\ T_2 \end{pmatrix}$$

that is obtained by solving a set of linear equations as:

$$\begin{pmatrix} O_x & O_y & 0 & 0 & 1 & 0 \\ 0 & 0 & O_x & O_y & 0 & 1 \\ P_{1x} & P_{1y} & 0 & 0 & 1 & 0 \\ 0 & 0 & P_{1x} & P_{1y} & 0 & 1 \\ P_{2x} & P_{2y} & 0 & 0 & 1 & 0 \\ 0 & 0 & P_{2x} & P_{2y} & 0 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \\ d \\ T_1 \\ T_2 \end{pmatrix} = \begin{pmatrix} O'_x \\ O'_y \\ P'_{1x} \\ P'_{1y} \\ P'_{2x} \\ P'_{2y} \end{pmatrix}$$

where (O_x, O_y) = O and x and y refer to the x and y coordinates of the points O, and so on.

Construction of Image Hash Table

[0031] Since occlusions, noise, and other changes can cause a triple of basis points on the given curve to not be visible in the corresponding image curve, affine coordinates of all points with respect to more sets of basis triple points may have to be recorded. The resulting Image Hash Table 5 (Figure 9) is a data structure that is a convenient way to represent this computed information so that the entries are the basis triples that give rise to a range of affine coordinates. The image hash table is constructed within the Hash Table Construction Module 4 using a suitable quantization of the affine coordinates, and recording the basis points that give rise to the respective affine coordinates. That is:

$$H(\alpha_1 \leq \alpha < \alpha_2, \beta_1 \leq \beta < \beta_2) = \{ \langle O', P'_1, P'_2 \rangle \dots \}$$

so that for any given affine coordinate (α, β) of a point, the possible basis points that gave rise to it can be found by looking in the hash table in the entry $\alpha_{\{1\}} \leq \alpha < \alpha_{\{2\}}$, $\beta_{\{1\}} \leq \beta < \beta_{\{2\}}$. Generalizing to the case of more curves in a curve group, the image hash table is constructed as follows. Each triple of consecutive features in a curve is used as a basis triple, and the affine coordinates of all features in the curve group are computed. Thus the basis points are

[0032] Because consecutive triples of features are used for basis points, only a linear number of basis points need to be recorded unlike $O(N^3)$ in straightforward geometric hashing. Also, the size of the hash table is $O(N^2)$ as against $O(N^4)$ in ordinary geometric hashing. The computational feasibility of this scheme together with its ability to localize objects makes it an improvement over existing variants of geometric hashing.

Indexing or word localization

[0033] Refer to the block diagram in Figure 10. During indexing, a Query Word 6 is given to the system, and curve groups are generated from the word using the pre-processing steps and requisite modules (7 and 8) for feature generation described in Figure 3. The word localization is attempted first using curve groups of longer average curve lengths. For each such curve group, sets of affine coordinates are computed within the Indexing Module 9 and used to index the Image Hash Table 12. Since the number of basis points are linear, this operation can be repeated with respect to all basis points in the curve group for robustness. For each basis triple that was indexed using the affine coordinates, the number of times it was indexed (called a hit) as well as the corresponding query triple are recorded. A histogram of the number of hits and the corresponding query word and matching basis points in the document image are recorded within the Histogram Ranking Module 10. The peaks in the histogram are then taken as the candidate locations for the query.

[0034] The indexing of the hash table accounts for the breaking of words into word segments in the image (or query word) by generating a set of affine coordinates as follows:

1. Let intra-word separation be: $T = (t_1, t_2)$.
2. For each basis triple $\langle O, P_1, P_2 \rangle$, and a given feature point P , compute affine coordinates (α, β) , and (α_k, β_k) where

$$\begin{bmatrix} \alpha'_k \\ \beta'_k \end{bmatrix} = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + \begin{bmatrix} (p_{1x} - o_x) & (p_{2x} - o_x) \\ (p_{1y} - o_y) & (p_{2y} - o_y) \end{bmatrix}^{-1} \begin{bmatrix} k_{n1} \\ k_{n2} \end{bmatrix}$$

- and where k is a number representative of the number of curves in a curve group. The value of k is meant to be tuned to the handwriting style of the author (i.e., the way he/she writes words in their characteristic style).
3. Use each of the affine coordinates to index the hash table and record peaks in the histogram of hits as described before.

Verification

[0035] The last step of word localization verifies the word at the candidate locations given in the indexing step. This is conducted by the Pose verification module 11. This step involves recovering the pose parameters (A, T) by solving the set of linear equations for the matching basis points corresponding to the significant hits.

[0036] Using the pose parameters, all points (i, j) (includes corner features) on curves of the query word are projected into the document image at location (i', j') where

$$\begin{pmatrix} i' \\ j' \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} i \\ j \end{pmatrix} + \begin{pmatrix} T_1 \\ T_2 \end{pmatrix}.$$

[0037] It is then verified if a point feature on each curve in the image lies within some neighborhood of the projected point. The ratio of matched projected points to the total number of points on all curves in the query word constitutes a verification score. The verification is said to succeed if this score is above a suitably chosen threshold. If no matching basis points are verified, then the next most significant query curve group is tried until no more significant groups are left. In practice, however, the correct query localization is achieved early in the indexing operation using the strongest

query curve.

Example

[0038] Figure 1 shows a scanned handwritten document and Figure 11A shows the result of pre-processing and feature extraction on that image. The corner features per curve used for hash table construction are shown as circles in Figure 11A. There are 179 curves and 2084 corners in all the curves combined. These give rise to 3494 basis points for the hash table. Figure 11B shows a query pattern consisting of a single curve. Figure 12 shows the histogram of hashing based on affine coordinates. Here the image basis points are plotted against the number of hits they obtained from affine coordinates on the query pattern. Figure 13 shows the results of hashing. The hashed image basis points corresponding to the three most significant peaks of the histogram are matched to their respective query basis triples to compute candidate poses. The query curve is then projected into the image using the pose parameters and shown overlayed on the original image in Figure 13. As can be seen, the top two matches localize the query pattern correctly at the two places it occurs. The third match is however, a false positive which can be removed during pose verification. The false positive occurs in this case because of a merging of the foreground text patterns with the lines of the tabular background in the image.

[0039] Referring back to Figures 2A-2C, illustration of query localization by hashing is shown, this time using curve groups. Figure 2A shows a sample document in which a word "database" occurs twice. The query word "database" is illustrated in Figure 2B. The inter-letter spacing between letters of the word is not uniform in the two instances. The query pattern used for indexing is shown in Figure 2C. Once again the top three matches are shown overlayed (after pose solution) on the original image to indicate query localization. Notice that using the indexing scheme, the word has been localized even when its constituent letters are written with different spacings in the two instances in which it occurs in the image. The false positive match shown here persisted even after pose verification, because of the similarity with the underlying word based on corner features.

Extension to Handwriting Tokenization

[0040] By choosing the query handwritten word to be one of the curve groups in the image itself, the above method can be used to identify multiple occurrences of the word in the document without explicitly matching to every single word in the document as is done by other tokenization schemes (e.g. DigiPaper™ by Xerox Corporation). Also, by using affine invariant features within curve groups, such a tokenization scheme is robust to changes in orientation, skew, and handwriting variances for a single author.

Generalizing to locating arbitrary 2d objects in scene images

[0041] By processing natural images to generate curves (perhaps by edge detection and curve tracing), the above method can be generalized to handle arbitrary 2d object shapes in unsegmented natural scene images. The grouping constraint to generate the curve groups may not be as easy to define in such cases as it was for handwritten documents (words are written more or less on a line). Finally, the above method admits other feature units besides corner features on curves. The grouping property, however, must be preserved with any feature unit used for localizing the object.

[0042] The foregoing description of the invention has been presented for purposes of illustration and to describe the best mode known for implementing of the invention. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obvious modifications or variations are possible in light of the above teachings. The embodiments were chosen and described in order to best illustrate the principles of the invention and its practical application to thereby enable one of ordinary skill in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated, as long as the principles described herein are followed. Thus, changes can be made in the above-described invention without departing from the intent and scope thereof. Therefore, it is intended that the specification and any examples be considered as exemplary only, with the true scope and spirit of the invention being indicated in the following claims.

Claims

1. A method of recognizing handwritten words in scanned documents, comprising:

processing a document containing handwriting wherein features for word localization are extracted from handwritten words contained in said document through basis points taken from a single curve of text lines, and wherein affine coordinates are computed for all features on all curves in a curve group;

storing said features in a memory;
 accessing said features from memory for comparison to handwritten words in a scanned document to recognize words within said scanned document.

5 2. The method of claim 1, wherein detection said text lines is independent of page orientation, and does not assume that the individual lines of handwritten text are parallel, and wherein said method does not require that word regions be aligned with text line orientation wherein intra-word statistics are derived from sample pages rather than using a fixed threshold.

10 3. A method of character recognition for handwritten words wherein features of handwritten word regions are extracted from said handwritten words for subsequent feature generation, extraction and recognition applications by:

generating connected regions for text lines of said handwritten words;
 conducting distance computation of said connected regions through along-line inter-region determination;
 15 determining curve and corner features from said connected regions; and
 developing intra-word text segment groupings.

4. The method of claim 3, wherein said text lines, and said text segments that lie along said text lines, are determined and intra-word separations are estimated wherein boundaries of a word segment region lying on said line are used to determine two extremal boundary points per region, and wherein curve extraction proceeds by determining said boundary points on connected component regions as points that have at least one "off" neighbor wherein said curves are next smoothed using a conventional line-segment approximation algorithm.

5. A system for the recognition of handwritten words wherein features of handwritten word regions are extracted from said handwritten words for subsequent feature generation, extraction and recognition applications, comprising a microprocessor programmed to:

generate connected regions for text lines of said handwritten words;
 conduct distance computations of said connected regions through along-line inter-region determination;
 30 determine curve and corner features from said connected regions; and
 develop intra-word text segment groupings;
 wherein said intra-word text segment groupings are the groupings of text segments into handwritten words that requires the knowledge of intra-word separations between text segments that lie along a line of text, text line detection is independent of page orientation and does not assume that the individual lines of handwritten text are parallel, and all word regions are not required to be aligned with the text line orientation.

6. A system according to claim 5, further comprising a memory wherein said handwritten words are indexed based on features for word localization extracted from said handwritten words contained in said document and wherein text segments are grouped from said handwritten words for purposes of indexing said documents based on words queries.

7. A system for recognizing handwritten words by pre-processing a scanned document containing handwriting where features for word localization are extracted from handwritten words contained in said document, said system comprising:

45 i) a feature extraction module which forms connecting component regions of a scanned document image representing said handwritten words of said document; and
 ii) a curve group generation module assembles groups of curves separated by intra-word separation curve segments belonging to said handwritten words within said scanned document image.

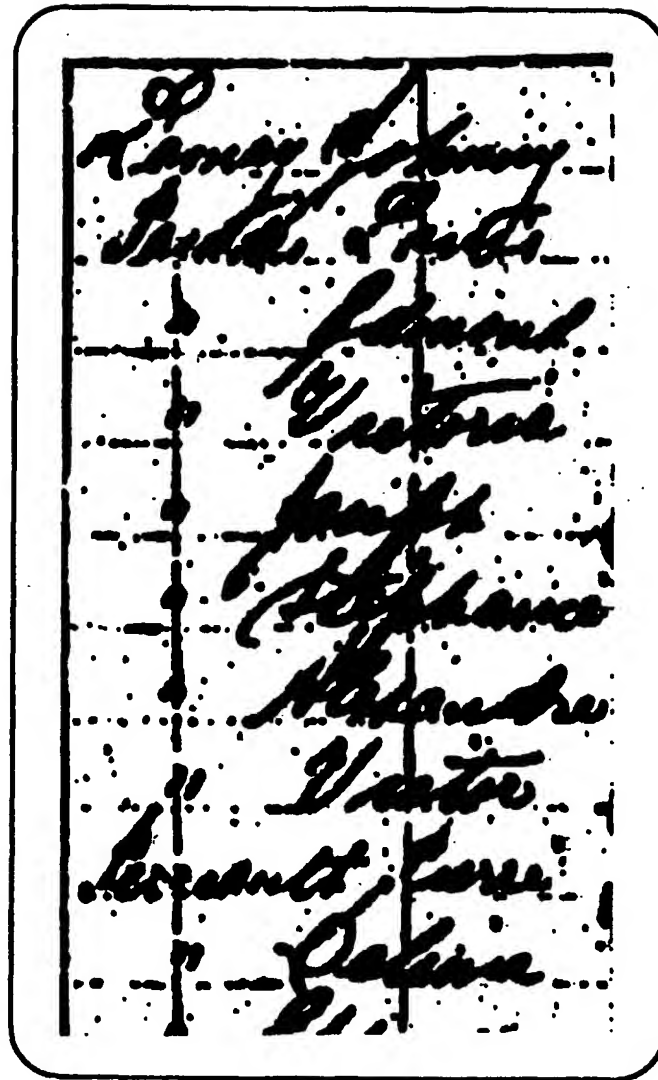


FIG. 1

While it is true that these problems exist also for image databases designed for a stand alone use, they are felt more keenly when such databases are designed cover)

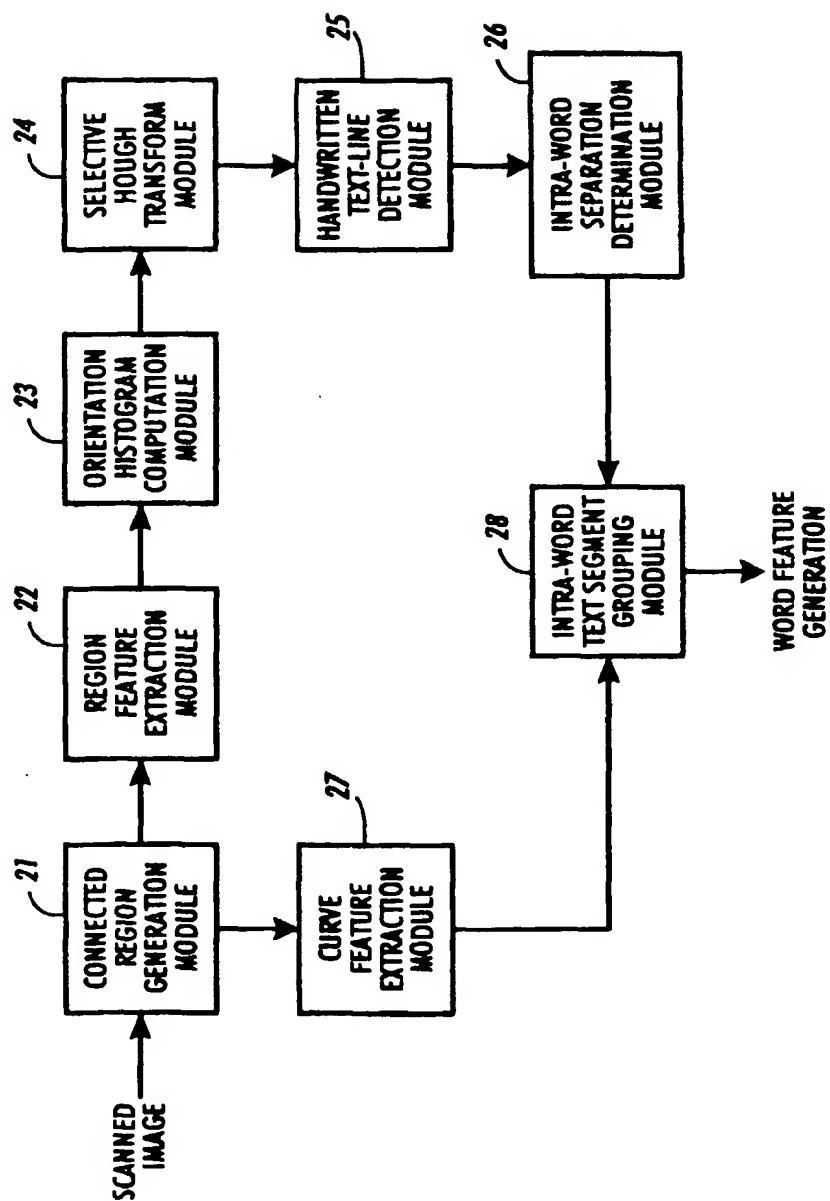
FIG. 2A


database

FIG. 2B

While it is true that these problems exist also for image ~~databases~~ designed for a stand ~~alone~~, they are felt more keenly when such ~~databases~~ are designed cover)

FIG. 2C

**FIG. 3**

**FIG. 4**

While it is true that there problems
exist also for image databases designed for
a stand alone use, they are felt more
keenly when such databases are designed
... ..

FIG. 5

~~It is true~~ it is true that these problems
 exist also for image databases designed for
 a stand alone use, they are for ~~the~~ ~~most~~
 many other such databases are designed
 over)

FIG. 6




FIG. 7

While it is true that these problems
 exist also for image data sets for
 a stand alone system, they are far more
 serious when data sets are distributed

FIG. 8




FIG. 9

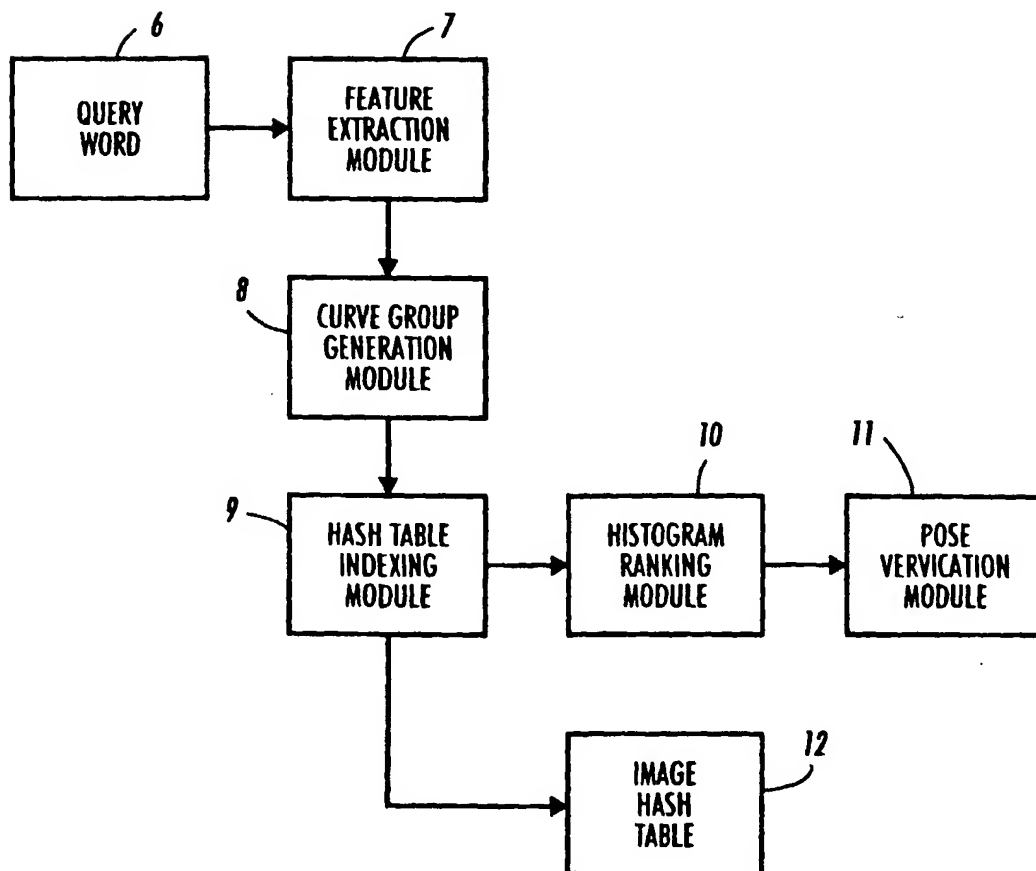


FIG. 10

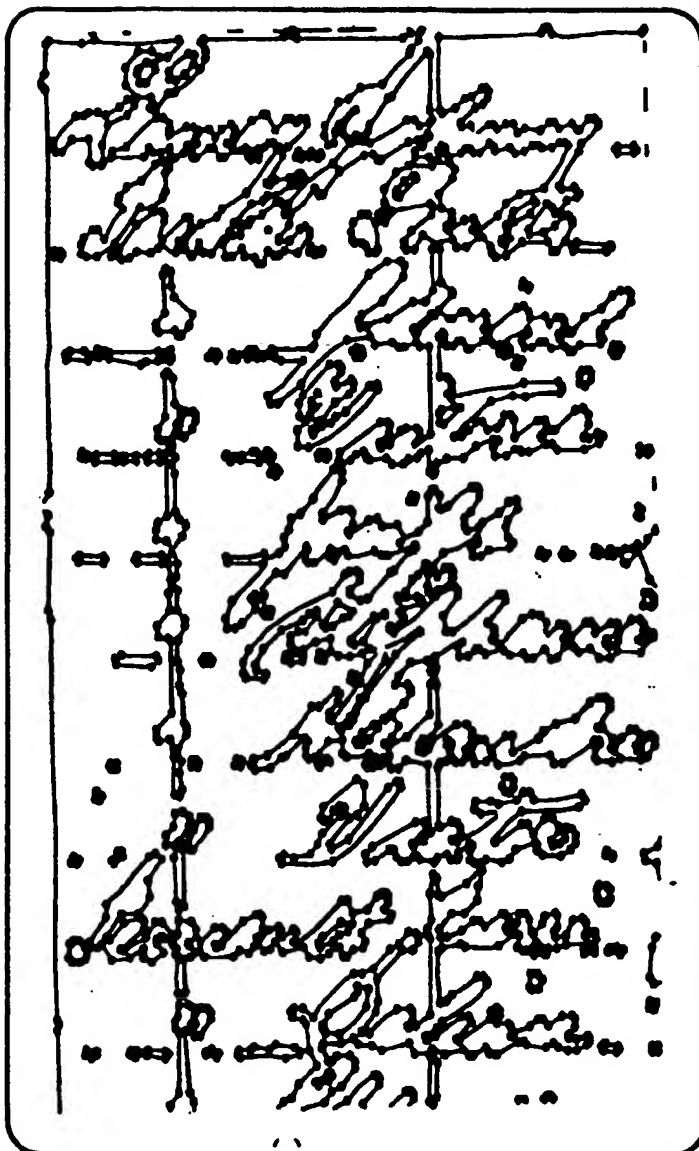


FIG. 11A




FIG. 11B




FIG. 12

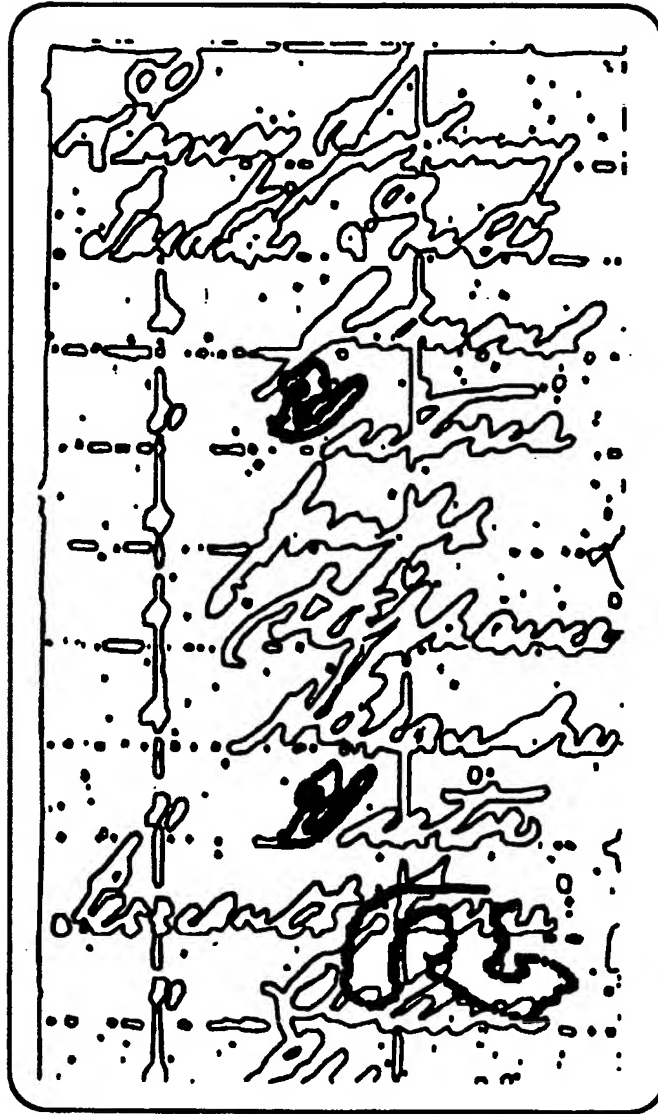


FIG. 13

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

☐ BLACK BORDERS

☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES

☐ FADED TEXT OR DRAWING

☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING

☐ SKEWED/SLANTED IMAGES

☒ COLOR OR BLACK AND WHITE PHOTOGRAPHS

☐ GRAY SCALE DOCUMENTS

☐ LINES OR MARKS ON ORIGINAL DOCUMENT

☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY

☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.